arXiv:1604.08697v1 [stat.ML] 29 Apr 2016

Sparse Generalized Eigenvalue Problem: Optimal Statistical Rates via Truncated Rayleigh Flow

Kean Ming Tan, Zhaoran Wang, Han Liu, and Tong Zhang

May 2, 2016

Abstract

Sparse generalized eigenvalue problem plays a pivotal role in a large family of high-dimensional learning tasks, including sparse Fisher's discriminant analysis, canonical correlation analysis, and sufficient dimension reduction. However, the theory of sparse generalized eigenvalue problem remains largely unexplored. In this paper, we exploit a non-convex optimization perspective to study this problem. In particular, we propose the truncated Rayleigh flow method (Rifle) to estimate the leading generalized eigenvector and show that it converges linearly to a solution with the optimal statistical rate of convergence. Our theory involves two key ingredients: (i) a new analysis of the gradient descent method on non-convex objective functions, as well as (ii) a fine-grained characterization of the evolution of sparsity patterns along the solution path. Thorough numerical studies are provided to back up our theory. Finally, we apply our proposed method in the context of sparse sufficient dimension reduction to two gene expression data sets.

1 Introduction

A broad class of high-dimensional statistical problems such as sparse canonical correlation analysis (CCA), sparse Fisher's discriminant analysis (FDA), and sparse sufficient dimension reduction (SDR) can be formulated as the sparse generalized eigenvalue problem (GEP). In detail, let $\mathbf{A} \in \mathbb{R}^{d\times d}$ be a symmetric matrix and $\mathbf{B} \in \mathbb{R}^{d\times d}$ be positive definite. For the symmetric-definite matrix pair $(\mathbf{A}, \mathbf{B})$, the sparse generalized eigenvalue problem aims to

obtain a sparse vector $\mathbf{v}^* \in \mathbb{R}^d$ satisfying

$$\mathbf{A}\mathbf{v}^* = \lambda_{\max}(\mathbf{A}, \mathbf{B}) \cdot \mathbf{B}\mathbf{v}^*, \tag{1}$$

where $\mathbf{v}^*$ is the leading generalized eigenvector corresponding to the largest generalized eigenvalue $\lambda_{\max}(\mathbf{A}, \mathbf{B})$ of the matrix pair $(\mathbf{A}, \mathbf{B})$. Let $s = \|\mathbf{v}^*\|_0$ be the number of non-zero entries in $\mathbf{v}^*$, and we assume that s is much smaller than d.

In real-world applications, the matrix pair $(\mathbf{A}, \mathbf{B})$ is a population quantity that is unknown in general. Instead, we can only access $(\widehat{\mathbf{A}}, \widehat{\mathbf{B}})$, which is an estimate of $(\mathbf{A}, \mathbf{B})$, i.e.,

$$\widehat{\mathbf{A}} = \mathbf{A} + \mathbf{E}_{\mathbf{A}} \quad \text{and} \quad \widehat{\mathbf{B}} = \mathbf{B} + \mathbf{E}_{\mathbf{B}},$$

where $\mathbf{E}_{\mathbf{A}}$ and $\mathbf{E}_{\mathbf{B}}$ are stochastic errors due to finite sample estimation. For those statistical applications considered in this paper, $\mathbf{E}_{\mathbf{A}}$ and $\mathbf{E}_{\mathbf{B}}$ are symmetric. We aim to approximate $\mathbf{v}^*$ based on $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ by approximately solving the following optimization problem

$$\underset{\mathbf{v} \in \mathbb{R}^d}{\operatorname{argmax}} \, \mathbf{v}^T \widehat{\mathbf{A}} \mathbf{v}, \quad \text{subject to } \mathbf{v}^T \widehat{\mathbf{B}} \mathbf{v} = 1, \quad \|\mathbf{v}\|_0 \leq s. \tag{2}$$

There are two major challenges in solving (2). Firstly, the problem in (2) requires maximizing a convex objective function over a non-convex set, which is NP-hard even if $\widehat{\mathbf{B}}$ is the identity matrix (Moghaddam et al., 2006a,b). Secondly, in the high-dimensional setting in which the dimension d is much larger than the sample size n, $\widehat{\mathbf{B}}$ is in general singular, and classical algorithms for solving generalized eigenvalue problems are not directly applicable (Golub and Van Loan, 2012).

In this paper, we propose a non-convex optimization algorithm to approximately solve (2). The proposed algorithm iteratively performs a gradient ascent step on the generalized Rayleigh quotient $\mathbf{v}^T \widehat{\mathbf{A}} \mathbf{v} / \mathbf{v}^T \widehat{\mathbf{B}} \mathbf{v}$, and a truncation step that preserves the top k entries of $\mathbf{v}$ with the largest magnitudes while setting the remaining entries to zero. Here k is a tuning parameter that controls the sparsity level of the solution. Strong theoretical guarantees are established for the proposed method. In particular, let $\{\mathbf{v}_t\}_{t=0}^L$ be the solution sequence resulting from the proposed algorithm, where L is the total number of iterations and $\mathbf{v}_0$ is the initialization point. We prove that, under mild conditions,

$$\|\mathbf{v}^t - \mathbf{v}^*\|_2 \leq \underbrace{\nu^t \cdot \|\mathbf{v}^0 - \mathbf{v}^*\|_2}_{\text{optimization error}} + \underbrace{\frac{\sqrt{\rho(\mathbf{E}_{\mathbf{A}}, 2k+s)^2 + \rho(\mathbf{E}_{\mathbf{B}}, 2k+s)^2}}{\xi(\mathbf{A}, \mathbf{B})}}_{\text{statistical error}} \quad (t = 1, \ldots, L). \tag{3}$$

The quantities $\nu \in (0, 1)$ and $\xi(\mathbf{A}, \mathbf{B})$ depend on the population matrix pair $(\mathbf{A}, \mathbf{B})$. These quantities will be specified in Section 4. Meanwhile, $\rho(\mathbf{E}_{\mathbf{A}}, 2k + s)$ is defined as

$$\rho(\mathbf{E}_{\mathbf{A}}, 2k + s) = \sup_{\|\mathbf{u}\|_2 = 1, \|\mathbf{u}\|_0 \leq 2k+s} |\mathbf{u}^T \mathbf{E}_{\mathbf{A}} \mathbf{u}| \tag{4}$$

and $\rho(\mathbf{E_B}, 2k+s)$ is defined similarly. In (3), the first term on the right-hand side quantifies the exponential decay of the optimization error, while the second term characterizes the statistical error due to finite sample estimation. In particular, for the aforementioned statistical problems such as sparse CCA, sparse FDA, and sparse SDR, we can show that

$$\max\{\rho(\mathbf{E_A}, 2k+s), \rho(\mathbf{E_B}, 2k+s)\} \leq \sqrt{\frac{(s+2k)\log d}{n}} \tag{5}$$

with high probability. Consequently, for any properly chosen k that is of the same order as s, the algorithm achieves an estimator of $\mathbf{v}^*$ with the optimal statistical rate of convergence $\sqrt{s \log d/n}$.

The sparse generalized eigenvalue problem in (2) is closely related to the classical matrix computation literature (see, e.g., Golub and Van Loan, 2012 for a survey, and more recent results in Ge et al., 2016). There are two key differences between our results and these previous works. First, we have an additional non-convex constraint on the sparsity level, which allows us to handle the high-dimensional setting. Second, due to the existence of stochastic errors, we allow the normalization matrix $\widehat{\mathbf{B}}$ to be rank-deficient, while in the classical setting $\widehat{\mathbf{B}}$ is assumed to be positive definite. In comparison with the existing generalized eigenvalue algorithms, our algorithm keeps the iterative solution sequence within a basin that only involves a few coordinates of $\mathbf{v}$ such that the corresponding submatrix of $\widehat{\mathbf{B}}$ is positive definite. Furthermore, in this way our algorithm ensures that the statistical errors in (3) only involves the largest sparse eigenvalues of the stochastic errors $\mathbf{E_A}$ and $\mathbf{E_B}$, which is defined in (4). In contrast, a straightforward application of the classical matrix perturbation theory gives a statistical error term that involves the largest eigenvalues of $\mathbf{E_A}$ and $\mathbf{E_B}$, which are much larger than their sparse eigenvalues (Stewart and Sun, 1990).

Notation: Let $\mathbf{v} = (v_1, \ldots, v_d)^T \in \mathbb{R}^d$. We define the ℓ_q-norm of $\mathbf{v}$ as $\|\mathbf{v}\|_q = (\sum_{j=1}^d |v_j|^q)^{1/q}$ for $1 \leq q < \infty$. Let $\lambda_{\max}(\mathbf{Z})$ and $\lambda_{\min}(\mathbf{Z})$ be the largest and smallest eigenvalues correspondingly. If $\mathbf{Z}$ is positive definite, we define its condition number as $\kappa(\mathbf{Z}) = \lambda_{\max}(\mathbf{Z})/\lambda_{\min}(\mathbf{Z})$. We denote $\lambda_k(\mathbf{Z})$ to be the kth eigenvalue of $\mathbf{Z}$, and the spectral norm of $\mathbf{Z}$ by $\|\mathbf{Z}\|_2 = \sup_{\|\mathbf{v}\|_2=1} \|\mathbf{Z}\mathbf{v}\|_2$. For $F \subset \{1, \ldots, d\}$, let $\mathbf{Z}_F \in \mathbb{R}^{s \times s}$ be the submatrix of $\mathbf{Z}$, where the rows and columns are restricted to the set F. We define $\rho(\mathbf{Z}, s) = \sup_{\|\mathbf{u}\|_2=1, \|\mathbf{u}\|_0 \leq s} |\mathbf{u}^T \mathbf{Z} \mathbf{u}|$.

2 Sparse Generalized Eigenvalue Problem and Its Applications

Many high-dimensional multivariate statistics methods can be formulated as special instances of (2). For instance, when $\widehat{\mathbf{B}} = \mathbf{I}$, (2) reduces to the sparse principal component analysis (PCA) that has received considerable attention within the past decade (among others, Zou et al., 2006; d'Aspremont et al., 2007, 2008; Witten et al., 2009; Ma, 2013; Cai et al., 2013; Yuan and Zhang, 2013; Vu et al., 2013; Vu and Lei, 2013; Birnbaum et al., 2013; Wang et al., 2014; Gu et al., 2014). In the following, we provide three examples when $\widehat{\mathbf{B}}$ is not the identity matrix. We start with sparse Fisher's discriminant analysis for classification problem (among others, Tibshirani et al., 2003; Guo et al., 2007; Leng, 2008; Clemmensen et al., 2012; Mai et al., 2012, 2015; Kolar and Liu, 2015; Gaynanova and Kolar, 2015; Fan et al., 2015).

Example 1. Sparse Fisher's discriminant analysis: Given n observations with K distinct classes, Fisher's discriminant problem seeks a low-dimensional projection of the observations such that the between-class variance, $\mathbf{\Sigma}_b$, is large relative to the within-class variance, $\mathbf{\Sigma}_w$. Let $\widehat{\mathbf{\Sigma}}_b$ and $\widehat{\mathbf{\Sigma}}_w$ be estimates of $\mathbf{\Sigma}_b$ and $\mathbf{\Sigma}_w$, respectively. To obtain a sparse leading discriminant vector, one solves

$$\underset{\mathbf{v}}{\text{maximize}}\ \mathbf{v}^T\widehat{\mathbf{\Sigma}}_b\mathbf{v}, \qquad \text{subject to } \mathbf{v}^T\widehat{\mathbf{\Sigma}}_w\mathbf{v} = 1, \qquad \|\mathbf{v}\|_0 \le s. \tag{6}$$

This is a special case of (2) with $\widehat{\mathbf{A}} = \widehat{\mathbf{\Sigma}}_b$ and $\widehat{\mathbf{B}} = \widehat{\mathbf{\Sigma}}_w$.

Next, we consider sparse canonical correlation analysis which explores the relationship between two high-dimensional random vectors (Witten et al., 2009; Chen et al., 2013; Gao et al., 2014, 2015).

Example 2. Sparse canonical correlation analysis: Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be two random vectors. Let $\mathbf{\Sigma}_x$ and $\mathbf{\Sigma}_y$ be the covariance matrices for $\boldsymbol{X}$ and $\boldsymbol{Y}$, respectively, and let $\mathbf{\Sigma}_{xy}$ be the cross-covariance matrix between $\boldsymbol{X}$ and $\boldsymbol{Y}$. To obtain sparse leading canonical direction vectors, we solve

$$\underset{\mathbf{v}_x,\mathbf{v}_y}{\text{maximize}}\ \mathbf{v}_x^T\widehat{\mathbf{\Sigma}}_{xy}\mathbf{v}_y, \qquad \text{subject to } \mathbf{v}_x^T\widehat{\mathbf{\Sigma}}_x\mathbf{v}_x = \mathbf{v}_y^T\widehat{\mathbf{\Sigma}}_y\mathbf{v}_y = 1, \quad \|\mathbf{v}_x\|_0 \le s_x, \quad \|\mathbf{v}_y\|_0 \le s_y, \tag{7}$$

where s_x and s_y control the cardinality of $\mathbf{v}_x$ and $\mathbf{v}_y$. This is a special case of (2) with

$$\widehat{\mathbf{A}} = \begin{pmatrix} \mathbf{0} & \widehat{\mathbf{\Sigma}}_{xy} \\ \widehat{\mathbf{\Sigma}}_{xy} & \mathbf{0} \end{pmatrix}, \qquad \widehat{\mathbf{B}} = \begin{pmatrix} \widehat{\mathbf{\Sigma}}_x & \mathbf{0} \\ \mathbf{0} & \widehat{\mathbf{\Sigma}}_y \end{pmatrix}, \qquad \mathbf{v} = \begin{pmatrix} \mathbf{v}_x \\ \mathbf{v}_y \end{pmatrix}.$$

Theoretical guarantees for sparse CCA were established recently. Chen et al. (2013) proposed a non-convex optimization algorithm for solving (7) with theoretical guarantees. However, their algorithm involves obtaining accurate estimators of $\mathbf{\Sigma}_x^{-1}$ and $\mathbf{\Sigma}_y^{-1}$, which are in general difficult to obtain without imposing additional structural assumptions on $\mathbf{\Sigma}_x^{-1}$ and $\mathbf{\Sigma}_y^{-1}$. In a follow-up work, Gao et al. (2014) proposed a two-stage procedure that attains the optimal statistical rate of convergence (Gao et al., 2015). However, they require the matrix $\mathbf{\Sigma}_{xy}$ to be low-rank, and the second stage of their procedure requires the normality assumption. In contrast, our proposal only requires the condition in (5), which is much weaker and holds for general sub-Gaussian distributions.

In the sequel, we consider a general regression problem with a univariate response Y and d-dimensional covariates $\boldsymbol{X}$, with the goal of inferring the conditional distribution of Y given $\boldsymbol{X}$. Sufficient dimension reduction is a popular approach for reducing the dimensionality of the covariates (Li, 1991; Cook and Lee, 1999; Cook, 2000, 2007; Cook and Forzani, 2008; Ma and Zhu, 2013). Many SDR problems can be formulated as generalized eigenvalue problems (Li, 2007; Chen et al., 2010). For simplicity, we consider a special case of SDR, the sparse sliced inverse regression.

Example 3. Sparse sliced inverse regression: Consider the model

$$Y = f(\mathbf{v}_1^T \boldsymbol{X}, \ldots, \mathbf{v}_K^T \boldsymbol{X}, \epsilon),$$

where ϵ is the stochastic error independent of $\boldsymbol{X}$, and $f(\cdot)$ is an unknown link function. Li (1991) proved that under certain regularity conditions, the subspace spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_K$ can be identified. Let $\mathbf{\Sigma}_x$ be the covariance matrix for $\boldsymbol{X}$ and let $\mathbf{\Sigma}_{E(\boldsymbol{X}|Y)}$ be the covariance matrix of the conditional expectation $E(\boldsymbol{X} \mid Y)$. The first leading eigenvector of the subspace spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_K$ can be identified by solving

$$\underset{\mathbf{v}}{\text{maximize}}\ \mathbf{v}^T \widehat{\mathbf{\Sigma}}_{E(\boldsymbol{X}|Y)} \mathbf{v}, \qquad \text{subject to } \mathbf{v}^T \widehat{\mathbf{\Sigma}}_x \mathbf{v} = 1, \qquad \|\mathbf{v}\|_0 \leq s. \tag{8}$$

This is a special case of (2) with $\widehat{\mathbf{A}} = \widehat{\mathbf{\Sigma}}_{E(\boldsymbol{X}|Y)}$ and $\widehat{\mathbf{B}} = \widehat{\mathbf{\Sigma}}_x$.

Many authors have proposed methodologies for sparse sliced inverse regression (Li and Nachtsheim, 2006; Zhu et al., 2006; Li and Yin, 2008; Chen et al., 2010; Yin and Hilafu, 2015). More generally, in the context of sparse SDR, Li (2007) and Chen et al. (2010) reformulated sparse SDR problems into the sparse generalized eigenvalue problem in (2). However, most of these approaches lack algorithmic and non-asymptotic statistical guarantees for the high-dimensional setting. Our results are directly applicable to many sparse SDR problems.

3 Truncated Rayleigh Flow Method

We propose an iterative algorithm to estimate $\mathbf{v}^*$, which we refer to as truncated Rayleigh flow method (Rifle). More specifically, the optimization problem (2) can be rewritten as

$$\underset{\mathbf{v}\in\mathbb{R}^d}{\text{maximize}} \frac{\mathbf{v}^T\widehat{\mathbf{A}}\mathbf{v}}{\mathbf{v}^T\widehat{\mathbf{B}}\mathbf{v}}, \qquad \text{subject to } \|\mathbf{v}\|_0 \le s,$$

where the objective function is called the generalized Rayleigh quotient. At each iteration of the algorithm, we compute the gradient of the generalized Rayleigh quotient and update the solution by its ascent direction. To achieve sparsity, a truncation operation is performed within each iteration. Let $\text{Truncate}(\mathbf{v}, F)$ be the truncated vector of $\mathbf{v}$ by setting $v_i = 0$ for $i \notin F$ for an index set F. We summarize the details in Algorithm 1.

Algorithm 1 Truncated Rayleigh Flow Method (Rifle)

Input: matrices $\widehat{\mathbf{A}}$, $\widehat{\mathbf{B}}$, initial vector $\mathbf{v}_0$, cardinality $k \in \{1, \ldots, d\}$, and step size η.
Let $t = 1$. Repeat the following until convergence:

1. $\rho_{t-1} \leftarrow \mathbf{v}_{t-1}^T\widehat{\mathbf{A}}\mathbf{v}_{t-1}/\mathbf{v}_{t-1}^T\widehat{\mathbf{B}}\mathbf{v}_{t-1}$.
2. $\mathbf{C} \leftarrow \mathbf{I} + (\eta/\rho_{t-1}) \cdot (\widehat{\mathbf{A}} - \rho_{t-1}\widehat{\mathbf{B}})$.
3. $\mathbf{v}_t' \leftarrow \mathbf{C}\mathbf{v}_{t-1}/\|\mathbf{C}\mathbf{v}_{t-1}\|_2$.
4. Let $F_t = \text{supp}(\mathbf{v}_t', k)$ contain the indices of $\mathbf{v}_t'$ with the largest k absolute values and $\text{Truncate}(\mathbf{v}_t', F_t)$ be the truncated vector of $\mathbf{v}_t'$ by setting $(\mathbf{v}_t')_i = 0$ for $i \notin F_t$.
5. $\widehat{\mathbf{v}}_t \leftarrow \text{Truncate}(\mathbf{v}_t', F_t)$.
6. $\mathbf{v}_t \leftarrow \widehat{\mathbf{v}}_t/\|\widehat{\mathbf{v}}_t\|_2$.
7. $t \leftarrow t + 1$.

Output: $\mathbf{v}_t$.

Algorithm 1 requires the choice of a step size η, a tuning parameter k on the cardinality of the solution, and an initialization $\mathbf{v}_0$. As suggested by our theory, we need η to be sufficiently small such that $\eta\lambda_{\max}(\widehat{\mathbf{B}}) < 1$. The tuning parameter k can be selected using cross-validation or based on prior knowledge. While the theoretical results in Theorem 1 suggests that a good initialization may be beneficial, our numerical results illustrate that the proposed algorithm is robust to different initialization $\mathbf{v}_0$. A more thorough discussion on this can be found in the following section.

4 Theoretical Results

In this section, we show that if the matrix pair $(\mathbf{A}, \mathbf{B})$ has a unique sparse leading eigenvector, then Algorithm 1 can accurately recover the population eigenvector from the noisy matrix pair $(\widehat{\mathbf{A}}, \widehat{\mathbf{B}})$. Recall from the introduction that $\mathbf{A}$ is symmetric and $\mathbf{B}$ is positive definite. This condition ensures that all generalized eigenvalues are real. We denote $\mathbf{v}^*$ to be the leading generalized eigenvector of $(\mathbf{A}, \mathbf{B})$. Let $V = \text{supp}(\mathbf{v}^*)$ be the index set corresponding to the non-zero elements of $\mathbf{v}^*$, and let $|V| = s$. Throughout the paper, for notational convenience, we employ λ_j and $\widehat{\lambda}_j$ to denote the jth generalized eigenvalue of the matrix pairs $(\mathbf{A}, \mathbf{B})$ and $(\widehat{\mathbf{A}}, \widehat{\mathbf{B}})$, respectively.

Our theoretical results depend on several quantities that are specific to the generalized eigenvalue problem. Let

$$\text{cr}(\mathbf{A}, \mathbf{B}) = \min_{\mathbf{v}:\|\mathbf{v}\|_2=1} \left[(\mathbf{v}^T\mathbf{A}\mathbf{v})^2 + (\mathbf{v}^T\mathbf{B}\mathbf{v})^2\right]^{1/2} > 0 \tag{9}$$

denote the Crawford number of the symmetric-definite matrix pair $(\mathbf{A}, \mathbf{B})$ (see, for instance, Stewart, 1979). For any set $F \subset \{1, \ldots, d\}$ with cardinality $|F| = k'$, let

$$\text{cr}(k') = \inf_{F:|F|\le k'} \text{cr}(\mathbf{A}_F, \mathbf{B}_F); \qquad \epsilon(k') = \sqrt{\rho(\mathbf{E}_\mathbf{A}, k')^2 + \rho(\mathbf{E}_\mathbf{B}, k')^2}, \tag{10}$$

where $\rho(\mathbf{E}_\mathbf{A}, k')$ is defined in (4). Moreover, let $\lambda_j = (F)$ and $\widehat{\lambda}_j(F)$ denote the jth generalized eigenvalue of the matrix pair $(\mathbf{A}_F, \mathbf{B}_F)$ and $(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$, respectively. In the following, we start with an assumption that these quantities are upper bounded for sufficiently large n.

Assumption 1. For any sufficiently large n, there exist constants $b, c > 0$ such that

$$\frac{\epsilon(k')}{\text{cr}(k')} \le b; \qquad \rho(\mathbf{E}_\mathbf{B}, k') \le c\lambda_{\min}(\mathbf{B})$$

for any $k' \ll n$, where $\epsilon(k')$ and $\text{cr}(k')$ are defined in (10).

Provided that n is large enough, it can be shown that the above assumption holds with high probability for most statistical models. More details can be found later in this section. We will use the following implications of Assumption 1 in the theoretical analysis, which are implied by matrix perturbation theory (Stewart, 1979; Stewart and Sun, 1990). In detail, by applications of Lemmas 1 and 2 in Appendix A, we have that for any $F \subset \{1, \ldots, d\}$ with $|F| = k'$, there exist constants a, c such that

$$(1-a)\lambda_j(F) \le \widehat{\lambda}_j(F) \le (1+a)\lambda_j(F); \qquad (1-c)\lambda_j(\mathbf{B}_F) \le \lambda_j(\widehat{\mathbf{B}}_F) \le (1+c)\lambda_j(\mathbf{B}_F),$$

and

$$c_{\text{lower}} \cdot \kappa(\mathbf{B}) \leq \kappa(\widehat{\mathbf{B}}_F) \leq c_{\text{upper}} \cdot \kappa(\mathbf{B}), \tag{11}$$

where $c_{\text{lower}} = (1-c)/(1+c)$ and $c_{\text{upper}} = (1+c)/(1-c)$. Here c is the same constant as in Assumption 1.

Our main results involve several more quantities. Let

$$\Delta\lambda = \min_{j>1} \frac{\lambda_1 - (1+a)\lambda_j}{\sqrt{1+\lambda_1^2}\sqrt{1+(1-a)^2\lambda_j^2}} \tag{12}$$

denote the eigengap for the generalized eigenvalue problem (Stewart, 1979; Stewart and Sun, 1990). Meanwhile, let

$$\gamma = \frac{(1+a)\lambda_2}{(1-a)\lambda_1}; \qquad \omega(k') = \frac{2}{\Delta\lambda \cdot \text{cr}(k')} \cdot \epsilon(k') \tag{13}$$

be an upper bound on the ratio between the second largest and largest generalized eigenvalues of the matrix pair $(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$, and the statistical error term, respectively. Let

$$\theta = 1 - \frac{1-\gamma}{30 \cdot (1+c) \cdot c_{\text{upper}}^2 \cdot \eta \cdot \lambda_{\max}(\mathbf{B}) \cdot \kappa^2(\mathbf{B}) \cdot [c_{\text{upper}}\kappa(\mathbf{B}) + \gamma]}$$

be some constant that depends on the matrix $\mathbf{B}$.

The following theorem shows that under suitable conditions, Algorithm 1 approximately recovers the leading generalized eigenvector $\mathbf{v}^*$.

Theorem 1. Let $k' = 2k + s$ and assume that $\Delta\lambda > \epsilon(k')/\text{cr}(k')$. Choose $k = Cs$ for sufficiently large C and η such that $\eta\lambda_{\max}(\mathbf{B}) < 1/(1+c)$ and

$$\nu = \sqrt{1 + 2[(s/k)^{1/2} + s/k]} \cdot \sqrt{1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma}\right]} < 1,$$

where c and c_{upper} are constants defined in Assumption 1 and (11). Given Assumption 1 and an initialization vector $\mathbf{v}_0$ with $\|\mathbf{v}_0\|_2 = 1$ satisfying

$$\frac{|(\mathbf{v}^*)^T\mathbf{v}_0|}{\|\mathbf{v}^*\|_2} \geq \omega(k') + \theta,$$

we have

$$\sqrt{1 - \frac{|(\mathbf{v}^*)^T\mathbf{v}_t|}{\|\mathbf{v}^*\|_2}} \leq \nu^t \cdot \sqrt{1 - \frac{|(\mathbf{v}^*)^T\mathbf{v}_0|}{\|\mathbf{v}^*\|_2}} + \sqrt{10} \cdot \frac{\omega(k')}{1-\nu}. \tag{14}$$

The proof of this theorem will be presented in Section 4.1. The implications of Theorem 1 are as follows. For the simplicity of discussion, assume that $(\mathbf{v}^*)^T\mathbf{v}_t$ and $(\mathbf{v}^*)^T\mathbf{v}_0$ are positive without loss of generality. Since $\mathbf{v}_t$ and $\mathbf{v}_0$ are unit vectors, from (14) we have

$$1 - \frac{|(\mathbf{v}^*)^T\mathbf{v}_t|}{\|\mathbf{v}^*\|_2} = \frac{1}{2}\left\|\mathbf{v}_t - \frac{\mathbf{v}^*}{\|\mathbf{v}^*\|_2}\right\|_2^2, \qquad 1 - \frac{|(\mathbf{v}^*)^T\mathbf{v}_0|}{\|\mathbf{v}^*\|_2} = \frac{1}{2}\left\|\mathbf{v}_0 - \frac{\mathbf{v}^*}{\|\mathbf{v}^*\|_2}\right\|_2^2.$$

In other words, (14) states that the ℓ_2 distance between $\mathbf{v}^*/\|\mathbf{v}^*\|_2$ and $\mathbf{v}_t$ can be upper bounded by two terms. The first term on the right-hand side of (14) quantifies the optimization error, which decreases to zero at a geometric rate since $\nu < 1$. Meanwhile, the second term on the right-hand side of (14) quantifies the statistical error $\omega(k')$. The quantity $\omega(k')$ depends on $\epsilon(k') = \sqrt{\rho(\mathbf{E_A}, k')^2 + \rho(\mathbf{E_B}, k')^2}$, where $\rho(\mathbf{E_A}, k')$ and $\rho(\mathbf{E_B}, k')$ are defined in (4). For a broad range of statistical models, these quantities converge to zero at the rate of $\sqrt{s \log d/n}$.

Theorem 1 involves a condition on initialization: the cosine of the angle between $\mathbf{v}^*$ and the initialization $\mathbf{v}_0$ needs to be strictly larger than a constant, since $\epsilon(k')$ goes to zero for sufficiently large n. There are several approaches to obtain such an initialization for different statistical models. For instance, Johnstone and Lu (2009) proposed a diagonal thresholding procedure to obtain an initial vector in the context of sparse PCA. Also, Chen et al. (2013) generalized the diagonal thresholding procedure for sparse CCA. An initial vector $\mathbf{v}_0$ can also be obtained by solving a convex relaxation of (2). This is explored in Gao et al. (2014) for sparse CCA. We refer the reader to Johnstone and Lu (2009) and Gao et al. (2014) for a more detailed discussion.

4.1 Proof Sketch of Theorem 1

To establish Theorem 1, we first quantify the error introduced by maximizing the empirical version of the generalized eigenvalue problem, restricted to a superset of V ($V \subset F$), that is,

$$\mathbf{v}(F) = \arg\max_{\mathbf{v}\in\mathbb{R}^d} \mathbf{v}^T\widehat{\mathbf{A}}\mathbf{v}, \qquad \text{subject to } \mathbf{v}^T\widehat{\mathbf{B}}\mathbf{v} = 1, \quad \text{supp}(\mathbf{v}) \subset F. \tag{15}$$

Then we establish an error bound between $\mathbf{v}_t'$ in Step 2 of Algorithm 1 and $\mathbf{v}(F)$. Finally, we quantify the error introduced by the truncated step in Algorithm 1. These results are stated in Lemmas 3–5 in the Appendix A.

4.2 Applications to Sparse CCA, PCA, and FDA

In this section, we provide some discussions on the implications of Theorem 1 in the context of sparse CCA, PCA, and FDA. First, we consider the sparse CCA model. For the sparse

CCA model, we assume

$$\begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{Y} \end{pmatrix} \sim N(\mathbf{0}, \boldsymbol{\Sigma}); \qquad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_x & \boldsymbol{\Sigma}_{xy} \\ \boldsymbol{\Sigma}_{xy}^T & \boldsymbol{\Sigma}_y \end{pmatrix}.$$

The following proposition characterizes the rate of convergence between $\widehat{\boldsymbol{\Sigma}}$ and $\boldsymbol{\Sigma}$. It follows from Lemma 6.5 of Gao et al. (2014).

Proposition 1. Let $\widehat{\boldsymbol{\Sigma}}_x$, $\widehat{\boldsymbol{\Sigma}}_y$, and $\widehat{\boldsymbol{\Sigma}}_{xy}$ be empirical estimates of $\boldsymbol{\Sigma}_x$, $\boldsymbol{\Sigma}_y$, and $\boldsymbol{\Sigma}_{xy}$, respectively. For any $C > 0$ and positive integer $\bar{k}$, there exists a constant $C' > 0$ such that

$$\rho(\widehat{\boldsymbol{\Sigma}}_x - \boldsymbol{\Sigma}_x, \bar{k}) \leq C\sqrt{\frac{\bar{k}\log d}{n}}; \quad \rho(\widehat{\boldsymbol{\Sigma}}_y - \boldsymbol{\Sigma}_y, \bar{k}) \leq C\sqrt{\frac{\bar{k}\log d}{n}}; \quad \rho(\widehat{\boldsymbol{\Sigma}}_{xy} - \boldsymbol{\Sigma}_{xy}, \bar{k}) \leq C\sqrt{\frac{\bar{k}\log d}{n}},$$

with probability greater than $1 - \exp(-C'\bar{k}\log d)$.

Recall from Example 2 the definitions of $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ in the context of sparse CCA. Choosing k to be of the same order as s, Proposition 1 implies that $\rho(\mathbf{E}_{\mathbf{A}}, k')$ and $\rho(\mathbf{E}_{\mathbf{B}}, k')$ in Theorem 1 are upper bounded by the order of $\sqrt{s\log d/n}$ with high probability. Thus, as the optimization error decays to zero, we obtain an estimator with the optimal statistical rate of convergence (Gao et al., 2015). In comparison with Chen et al. (2013); Gao et al. (2014), we do not require structural assumptions on $\boldsymbol{\Sigma}_x$, $\boldsymbol{\Sigma}_y$, and $\boldsymbol{\Sigma}_{xy}$, as discussed in Section 2. Furthermore, we can establish the same error bounds as in Proposition 1 for sub-Gaussian distributions. Since our theory only relies on these error bounds, rather than the distributions of $\boldsymbol{X}$ and $\boldsymbol{Y}$, we can easily handle sub-Gaussian distributions in comparison with the results in Gao et al. (2014).

In addition, our results have direct implications on sparse PCA and sparse FDA. For the sparse PCA, assume that $\boldsymbol{X} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$. As mentioned in Section 2, sparse PCA is a special case of sparse generalized eigenvalue problem when $\widehat{\mathbf{B}} = \mathbf{I}$ and $\widehat{\mathbf{A}} = \widehat{\boldsymbol{\Sigma}}$. Theorem 1 and Proposition 1 imply that our estimator achieves the optimal statistical rate of convergence of $\sqrt{s\log d/n}$ (Cai et al., 2013; Vu and Lei, 2013). Recall the sparse FDA problem in Example 1. For simplicity, we assume there are two classes with mean $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$, and covariance matrix $\boldsymbol{\Sigma}$. Then it can be shown that the between-class and within-class covariance matrices take the form $\boldsymbol{\Sigma}_b = (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T$ and $\boldsymbol{\Sigma}_w = \boldsymbol{\Sigma}$. If we further assume the data are sub-Gaussian, using results similar to Proposition 1, we have $\rho(\widehat{\boldsymbol{\Sigma}}_b - \boldsymbol{\Sigma}_b, k')$ and $\rho(\widehat{\boldsymbol{\Sigma}}_w - \boldsymbol{\Sigma}_w, k')$ are upper bounded by the order of $\sqrt{s\log d/n}$ with high probability. Similar results were established in Fan et al. (2015).

5 Numerical Studies

We perform numerical studies to evaluate the performance of our proposal, Rifle, compared to some existing methods. We consider sparse Fisher's discriminant analysis and sparse canonical correlation analysis, each of which can be recast as the sparse generalized eigenvalue problem (2), as shown in Examples 1–2. Our proposal involves an initial vector $\mathbf{v}_0$ and a tuning parameter k on the cardinality. In our simulation studies, we generate each entry of $\mathbf{v}_0$ from a standard normal distribution. We then standardize the vector $\mathbf{v}_0$ such that $\|\mathbf{v}_0\|_2 = 1$. We first run Algorithm 1 with a large value of truncation parameter k. The solution is then used as the initial value for Algorithm 1 with a smaller value of k. This type of initialization has been considered in the context of sparse PCA and is shown to yield good empirical performance (Yuan and Zhang, 2013).

5.1 Fisher's Discriminant Analysis

We consider high-dimensional classification problem using sparse Fisher's discriminant analysis. The data consists of an $n \times d$ matrix $\mathbf{X}$ with d features measured on n observations, each of which belongs to one of K classes. We let $\mathbf{x}_i$ denote the ith row of $\mathbf{X}$, and let $C_k \subset \{1, \ldots, n\}$ contains the indices of the observations in the kth class with $n_k = |C_k|$ and $\sum_{k=1}^K n_k = n$.

Recall from Example 1 that this is a special case of the sparse generalized eigenvalue problem with $\widehat{\mathbf{A}} = \widehat{\mathbf{\Sigma}}_b$ and $\widehat{\mathbf{B}} = \widehat{\mathbf{\Sigma}}_w$. Let $\widehat{\boldsymbol{\mu}}_k = \sum_{i \in C_k} \mathbf{x}_i / n_k$ be the estimated mean for the kth class. The standard estimates for $\mathbf{\Sigma}_w$ and $\mathbf{\Sigma}_b$ are

$$\widehat{\mathbf{\Sigma}}_w = \frac{1}{n} \sum_{k=1}^K \sum_{i \in C_k} (\mathbf{x}_i - \widehat{\boldsymbol{\mu}}_k)(\mathbf{x}_i - \widehat{\boldsymbol{\mu}}_k)^T; \qquad \widehat{\mathbf{\Sigma}}_b = \frac{1}{n} \sum_{k=1}^K n_k \widehat{\boldsymbol{\mu}}_k \widehat{\boldsymbol{\mu}}_k^T.$$

We consider two simulation settings similar to that of Witten et al. (2009):

1. Binary classification: in this example, we set $\boldsymbol{\mu}_1 = \mathbf{0}$, $\mu_{2j} = 0.5$ for $j = \{2, 4, \ldots, 40\}$, and $\mu_{2j} = 0$ otherwise. Let $\mathbf{\Sigma}$ be a block diagonal covariance matrix with five blocks, each of dimension $d/5 \times d/5$. The (j, j')th element of each block takes value $0.7^{|j-j'|}$. As suggested by Witten et al. (2009), this covariance structure is intended to mimic the covariance structure of gene expression data. The data are simulated as $\mathbf{x}_i \sim N(\boldsymbol{\mu}_k, \mathbf{\Sigma})$ for $i \in C_k$.

2. Multi-class classification: there are $K = 4$ four classes in this example. Let $\mu_{kj} = (k-1)/3$ for $j = \{2, 4, \ldots, 40\}$ and $\mu_{kj} = 0$ otherwise. The data are simulated as $\mathbf{x}_i \sim N(\boldsymbol{\mu}_k, \mathbf{\Sigma})$ for $i \in C_k$, with the same covariance structure for binary classification.

As noted in Witten et al. (2009), a one-dimensional vector projection of the data fully captures the class structure.

Four approaches are compared: (i) our proposal Rifle; (ii) ℓ_1-penalized logistic or multinomial regression implemented using the R package glmnet; (iii) ℓ_1-penalized FDA with diagonal estimate of $\mathbf{\Sigma}_w$ implemented using the R package penalizedLDA (Witten et al., 2009); and (iv) direct approach to sparse discriminant analysis (Mai et al., 2012, 2015) implemented using the R package dsda and msda for binary and multi-class classification, respectively. For each method, models are fit on the training set with tuning parameter selected using 5-fold cross-validation. Then, the models are evaluated on the test set. In addition to the aforementioned models, we consider an oracle estimator using the theoretical direction $\mathbf{v}^*$, computed using the population quantities $\mathbf{\Sigma}_w$ and $\mathbf{\Sigma}_b$.

To compare the performance of the different proposals, we report the misclassification error on the test set and the number of non-zero features selected in the models. The results for 500 training samples and 1000 test samples, with $d = 1000$ features, are reported in Table 1. From Table 1, we see that our proposal have the lowest misclassification error compared to other competing methods. Witten et al. (2009) has the highest misclassification error in both of our simulation settings, since it does not take into account the dependencies among the features. Mai et al. (2012) and Mai et al. (2015) perform slightly worse than our proposal in terms of misclassification error. Moreover, they use a large number of features in their model, which renders interpretation difficult. In contrast, the number of features selected by our proposal is very close to that of the oracle estimator.

Table 1: The number of misclassified observations out of 1000 test samples and number of non-zero features (and standard errors) for binary and multi-class classification problems, averaged over 200 data sets. The results are for models trained with 500 training samples with 1000 features.

		ℓ_1-penalized	ℓ_1-FDA	direct	our proposal	oracle
Binary	Error	29.21 (0.59)	295.85 (1.08)	26.92 (0.49)	19.18 (0.56)	8.23 (0.19)
	Features	111.79 (1.50)	23.45 (0.82)	122.85 (2.17)	49.15 (0.57)	41 (0)
Multi-class	Error	497.43 (1.69)	495.61 (1.14)	244.25 (1.39)	213.82 (1.56)	153.18 (0.80)
	Features	61.99 (2.51)	23.76 (2.23)	109.30 (2.33)	55.68 (0.36)	41 (0)

5.2 Canonical Correlation Analysis

In this section, we study the relationship between two sets of random variables $\boldsymbol{X} \in \mathbb{R}^{d/2}$ and $\boldsymbol{Y} \in \mathbb{R}^{d/2}$ in the high-dimensional setting using sparse CCA. Let $\boldsymbol{\Sigma}_x$, $\boldsymbol{\Sigma}_y$, and $\boldsymbol{\Sigma}_{xy}$ be the covariance matrices of $\boldsymbol{X}$ and $\boldsymbol{Y}$, and cross-covariance matrix of $\boldsymbol{X}$ and $\boldsymbol{Y}$, respectively. Assume that $(\boldsymbol{X}, \boldsymbol{Y}) \sim N(\mathbf{0}, \boldsymbol{\Sigma})$ with

$$\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_x & \boldsymbol{\Sigma}_{xy} \\ \boldsymbol{\Sigma}_{xy} & \boldsymbol{\Sigma}_y \end{pmatrix}; \qquad \boldsymbol{\Sigma}_{xy} = \boldsymbol{\Sigma}_x \mathbf{v}_x^* \lambda_1 (\mathbf{v}_y^*)^T \boldsymbol{\Sigma}_y,$$

where $0 < \lambda_1 < 1$ is the largest generalized eigenvalue and $\mathbf{v}_x^*$ and $\mathbf{v}_y^*$ are the leading pair of canonical directions. The data consists of two $n \times (d/2)$ matrices $\mathbf{X}$ and $\mathbf{Y}$. We assume that each row of the two matrices are generated according to $(\mathbf{x}_i, \mathbf{y}_i) \sim N(\mathbf{0}, \boldsymbol{\Sigma})$. The goal of CCA is to estimate the canonical directions $\mathbf{v}_x^*$ and $\mathbf{v}_y^*$ based on the data matrices $\mathbf{X}$ and $\mathbf{Y}$.

Let $\widehat{\boldsymbol{\Sigma}}_x$, $\widehat{\boldsymbol{\Sigma}}_y$ be the sample covariance matrices of $\boldsymbol{X}$ and $\boldsymbol{Y}$, and let $\widehat{\boldsymbol{\Sigma}}_{xy}$ be the sample cross-covariance matrix of $\boldsymbol{X}$ and $\boldsymbol{Y}$. Recall from Example 2 that the sparse CCA problem can be recast as the generalized eigenvalue problem with

$$\widehat{\mathbf{A}} = \begin{pmatrix} \mathbf{0} & \widehat{\boldsymbol{\Sigma}}_{xy} \\ \widehat{\boldsymbol{\Sigma}}_{xy} & \mathbf{0} \end{pmatrix}, \qquad \widehat{\mathbf{B}} = \begin{pmatrix} \widehat{\boldsymbol{\Sigma}}_x & \mathbf{0} \\ \mathbf{0} & \widehat{\boldsymbol{\Sigma}}_y \end{pmatrix}, \qquad \mathbf{v} = \begin{pmatrix} \mathbf{v}_x \\ \mathbf{v}_y \end{pmatrix}.$$

In our simulation setting, we set $\lambda_1 = 0.9$, $v_{x,j}^* = v_{y,j}^* = 1/\sqrt{5}$ for $j = \{1, 6, 11, 16, 21\}$, and $v_{x,j}^* = v_{y,j}^* = 0$ otherwise. Then, we normalize $\mathbf{v}_x^*$ and $\mathbf{v}_y^*$ such that $(\mathbf{v}_x^*)^T \boldsymbol{\Sigma}_x \mathbf{v}_x^* = (\mathbf{v}_y^*)^T \boldsymbol{\Sigma}_y \mathbf{v}_y^* = 1$. We consider the case when $\boldsymbol{\Sigma}_x$ and $\boldsymbol{\Sigma}_y$ are block diagonal matrix with five blocks, each of dimension $d/5 \times d/5$, where the (j, j')th element of each block takes value $0.7^{|j-j'|}$.

We compare our proposal to Witten et al. (2009), implemented using the R package PMA. Their proposal involves choosing two tuning parameters that controls the sparsity of the estimated directional vectors, which we select using cross-validation. We perform our method using multiple values of k, and report results for $k = \{15, 25\}$ since they are qualitatively similar for large n. The output of both our proposal and that of Witten et al. (2009) are normalized to have norm one, whereas the true parameters $\mathbf{v}_x^*$ and $\mathbf{v}_y^*$ are normalized with respect to $\boldsymbol{\Sigma}_x$ and $\boldsymbol{\Sigma}_y$. To evaluate the performance of the two methods, we normalize $\mathbf{v}_x^*$ and $\mathbf{v}_y^*$ such that they have norm one, and compute the squared ℓ_2 distance between the estimated and the true directional vectors. The results for $d = 600$, $s = 10$, averaged over 100 data sets, are plotted in Figure 1.

From Figure 1, we see that our proposal outperforms Witten et al. (2009) uniformly across different sample sizes. This is not surprising since Witten et al. (2009) uses diagonal estimates of $\boldsymbol{\Sigma}_x$ and $\boldsymbol{\Sigma}_y$ to compute the directional vectors. Our proposal is not sensitive to

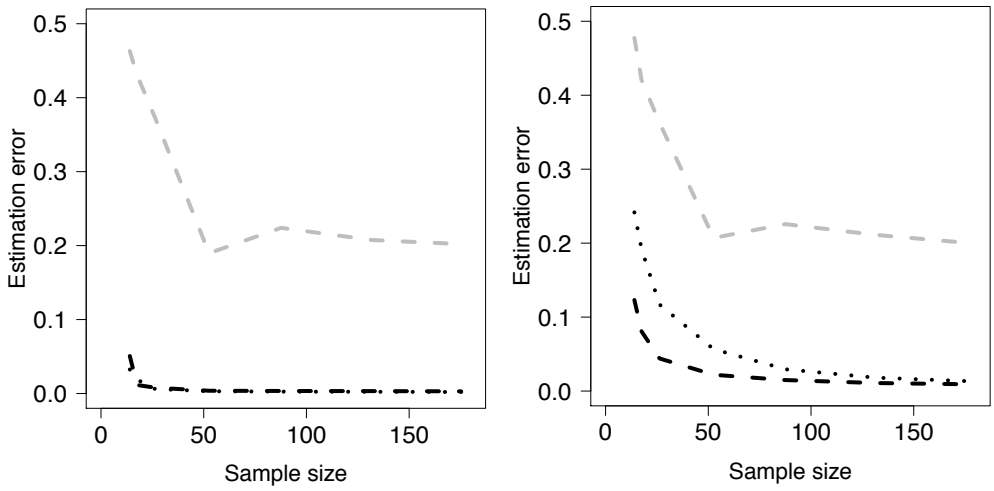

(a) (b)

Figure 1: The squared ℓ_2 distance plotted as a function of scaled sample size $n/\{s\log(d)\}$ for $d = 600$, $s = 10$. Panels (a) and (b) are results for the left and right canonical directions, respectively. The different curves represent our proposal with $k = 15$ (long-dashed), our proposal with $k = 25$ (dots), and Witten et al. (2009) (grey long-dashed).

the truncation parameter k when the sample size is sufficiently large. In addition, we see that the squared ℓ_2 distance for our proposal is inversely proportional to $n/\{s\log(d)\}$, as suggested by Theorem 1.

6 Data Application

In this section, we apply our method in the context of sparse sliced inverse regression as in Example 3. The data sets we consider are:

1. Leukemia (Golub et al., 1999): 7,129 gene expression measurements from 25 patients with acute myeloid leukemia and 47 patients with acute lymphoblastic luekemia. The data are available from `http://www.broadinstitute.org/cgi-bin/cancer/datasets.cgi`. Recently, this data set is analyzed in the context of sparse sufficient dimension reduction in Yin and Hilafu (2015).

2. Lung cancer (Spira et al., 2007): 22,283 gene expression measurements from large airway epithelial cells sampled from 97 smokers with lung cancer and 90 smokers without lung cancer. The data are publicly available from GEO at accession number GDS2771.

We preprocess the leukemia data set following Golub et al. (1999) and Yin and Hilafu (2015). In particular, we set gene expression readings of 100 or fewer to 100, and expression readings of 16,000 or more to 16,000. We then remove genes with difference and ratio between the maximum and minimum readings that are less than 500 and 5, respectively. A log-transformation is then applied to the data. This gives us a data matrix $\mathbf{X}$ with 72 rows/samples and 3571 columns/genes. For the lung cancer data, we simply select the 2,000 genes with the largest variance as in Petersen et al. (2015). This gives a data matrix with 167 rows/samples and 2,000 columns/genes. We further standardize both the data sets so that the genes have mean equals zero and variance equals one.

Recall from Example 3 that in order to apply our method, we need the estimates $\widehat{\mathbf{A}} = \widehat{\mathbf{\Sigma}}_{E(\boldsymbol{X}|Y)}$ and $\widehat{\mathbf{B}} = \widehat{\mathbf{\Sigma}}_x$. The quantity $\widehat{\mathbf{\Sigma}}_x$ is simply the sample covariance matrix of $\boldsymbol{X}$. Let n_1 and n_2 be the number of samples of the two classes in the data set. Let $\widehat{\mathbf{\Sigma}}_{x,1}$ and $\widehat{\mathbf{\Sigma}}_{x,2}$ be the sample covariance matrix calculated using only data from class one and class two, respectively. Then, the covariance matrix of the conditional expectation can be estimated by

$$\widehat{\mathbf{\Sigma}}_{E[\boldsymbol{X}|Y]} = \widehat{\mathbf{\Sigma}}_x - \frac{1}{n}\sum_{k=1}^{2} n_k \widehat{\mathbf{\Sigma}}_{x,k},$$

where $n = n_1 + n_2$ (Li, 1991; Li and Nachtsheim, 2006; Zhu et al., 2006; Li and Yin, 2008; Chen et al., 2010; Yin and Hilafu, 2015). As mentioned in Section 5, we run Algorithm 1 with a large value of k, and used its solution as the initial value for Algorithm 1 with a smaller value of k. Let $\widehat{\mathbf{v}}_t$ be the output of Algorithm 1. Similar to Yin and Hilafu (2015), we plot the box-plot of the sufficient predictor, $\mathbf{X}\widehat{\mathbf{v}}_t$, for the two classes in each data set. The results with $k = 25$ for leukemia and lung cancer data sets are in Figures 2(a)-(b), respectively.

From Figure 2(a), for the leukemia data set, we see that the sufficient predictor for the two groups are much more well separated than the results in Yin and Hilafu (2015). Moreover, our proposal is a one-step procedure with theoretical guarantees whereas their proposal is sequential without theoretical guarantees. For the lung cancer data set, we see that there is some overlap between the sufficient predictor for subjects with and without lung cancer. These results are consistent in the literature where it is known that the lung cancer data set is a much more difficult classification problem compared to that of the leukemia data set (Fan and Fan, 2008; Petersen et al., 2015).

7 Discussion

We propose the truncated Rayleigh flow for solving sparse generalized eigenvalue problem. The proposed method successfully handles ill-conditioned normalization matrices that arise

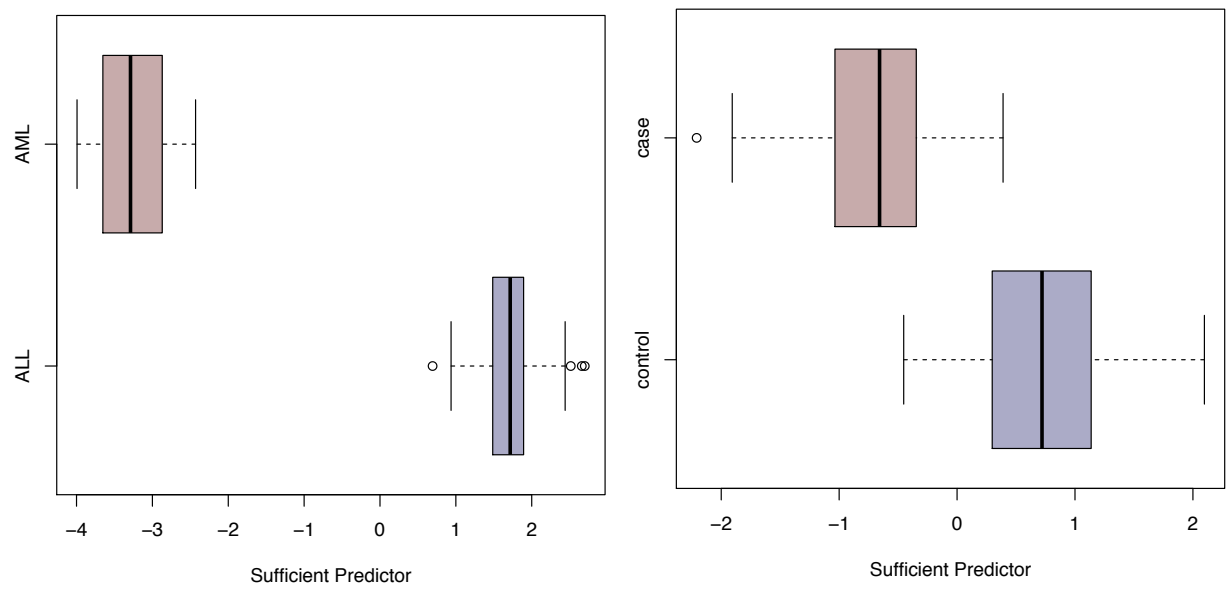

(a) (b)

Figure 2: Panels (a) and (b) contain box-plots of the sufficient predictor $\mathbf{X}\widehat{\mathbf{v}}_t$ obtained from Algorithm 1 for the leukemia and lung cancer data sets. In panel (a), the y-axis represents patients with acute lymphoblastic leukemia (ALL) and acute myeloid leukemia (AML), respectively. In panel (b), the y-axis represents patients with and without lung cancer, respectively.

from the high-dimensional setting due to finite sample estimation, and enjoys geometric convergence to a solution with the optimal statistical rate of convergence. The proposed method and theory have applications to a broad family of important statistical problems including sparse FDA, sparse CCA, and sparse SDR. Compared to existing theory, our theory does not require any structural assumption on $(\mathbf{A}, \mathbf{B})$, nor normality assumption on the data.

Acknowledgement

We thank Ashley Petersen for providing us the lung cancer data set. Han Liu was supported by NSF CAREER Award DMS1454377, NSF IIS1408910, NSF IIS1332109, NIH R01MH102339, NIH R01GM083084, and NIH R01HG06841. Tong Zhang was supported by NSF IIS-1250985, NSF IIS-1407939, and NIH R01AI116744. Kean Ming Tan was supported by NSF IIS-1250985 and NSF IIS-1407939. Zhaoran Wang was supported by Microsoft Research PhD fellowship.

References

Birnbaum, A., Johnstone, I. M., Nadler, B. and Paul, D. (2013). Minimax bounds for sparse PCA with noisy high-dimensional data. *Annals of Statistics* **41** 1055–1084.

Cai, T. T., Ma, Z. and Wu, Y. (2013). Sparse PCA: Optimal rates and adaptive estimation. *Annals of Statistics* **41** 3074–3110.

Chen, M., Gao, C., Ren, Z. and Zhou, H. H. (2013). Sparse CCA via precision adjusted iterative thresholding. *arXiv preprint arXiv:1311.6186* .

Chen, X., Zou, C. and Cook, R. D. (2010). Coordinate-independent sparse sufficient dimension reduction and variable selection. *Annals of Statistics* **38** 3696–3723.

Clemmensen, L., Hastie, T., Witten, D. and Ersbøll, B. (2012). Sparse discriminant analysis. *Technometrics* **53** 406–413.

Cook, R. D. (2000). SAVE: a method for dimension reduction and graphics in regression. *Communications in Statistics - Theory and Methods* **29** 2109–2121.

Cook, R. D. (2007). Fisher lecture: Dimension reduction in regression. *Statistical Science* **22** 1–26.

Cook, R. D. and Forzani, L. (2008). Principal fitted components for dimension reduction in regression. *Statistical Science* **23** 485–501.

Cook, R. D. and Lee, H. (1999). Dimension reduction in binary response regression. *Journal of the American Statistical Association* **94** 1187–1200.

d'Aspremont, A., Bach, F. and Ghaoui, L. E. (2008). Optimal solutions for sparse principal component analysis. *The Journal of Machine Learning Research* **9** 1269–1294.

d'Aspremont, A., El Ghaoui, L., Jordan, M. I. and Lanckriet, G. R. (2007). A direct formulation for sparse PCA using semidefinite programming. *SIAM Review* **49** 434–448.

Fan, J. and Fan, Y. (2008). High dimensional classification using features annealed independence rules. *Annals of statistics* **36** 2605–2637.

Fan, J., Ke, Z. T., Liu, H. and Xia, L. (2015). QUADRO: A supervised dimension reduction method via Rayleigh quotient optimization. *Annals of Statistics* **43** 1498.

Gao, C., Ma, Z., Ren, Z. and Zhou, H. H. (2015). Minimax estimation in sparse canonical correlation analysis. *Annals of Statistics* **43** 2168–2197.

Gao, C., Ma, Z. and Zhou, H. H. (2014). Sparse CCA: Adaptive estimation and computational barriers. *arXiv preprint arXiv:1409.8565* .

Gaynanova, I. and Kolar, M. (2015). Optimal variable selection in multi-group sparse discriminant analysis. *Electronic Journal of Statistics* **9** 2007–2034.

Ge, R., Jin, C., Kakade, S. M., Netrapalli, P. and Sidford, A. (2016). Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis. *arXiv preprint arXiv:1604.03930* .

Golub, G. H. and Van Loan, C. F. (2012). *Matrix Computations*, vol. 3. JHU Press.

Golub, T. R., Slonim, D. K., Tamayo, P. et al. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *science* **286** 531–537.

Gu, Q., Wang, Z. and Liu, H. (2014). Sparse PCA with oracle property. In *Advances in Neural Information Processing Systems*.

Guo, Y., Hastie, T. and Tibshirani, R. (2007). Regularized linear discriminant analysis and its application in microarrays. *Biostatistics* **8** 86–100.

Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104** 682–693.

Kolar, M. and Liu, H. (2015). Optimal feature selection in high-dimensional discriminant analysis. *IEEE Transactions on Information Theory* **61** 1063–1083.

Leng, C. (2008). Sparse optimal scoring for multiclass cancer diagnosis and biomarker detection using microarray data. *Computational Biology and Chemistry* **32** 417–425.

Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86** 316–327.

Li, L. (2007). Sparse sufficient dimension reduction. *Biometrika* **94** 603–613.

Li, L. and Nachtsheim, C. J. (2006). Sparse sliced inverse regression. *Technometrics* **48** 503–510.

Li, L. and Yin, X. (2008). Sliced inverse regression with regularizations. *Biometrics* **64** 124–131.

Ma, Y. and Zhu, L. (2013). A review on dimension reduction. *International Statistical Review* **81** 134–150.

Ma, Z. (2013). Sparse principal component analysis and iterative thresholding. *Annals of Statistics* **41** 772–801.

Mai, Q., Yang, Y. and Zou, H. (2015). Multiclass sparse discriminant analysis. *arXiv preprint arXiv:1504.05845* .

Mai, Q., Zou, H. and Yuan, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika* **99** 29–42.

Moghaddam, B., Weiss, Y. and Avidan, S. (2006a). Generalized spectral bounds for sparse LDA. In *International Conference on Machine Learning*.

Moghaddam, B., Weiss, Y. and Avidan, S. (2006b). Spectral bounds for sparse PCA: Exact and greedy algorithms. In *Advances in Neural Information Processing Systems*.

Petersen, A., Witten, D. and Simon, N. (2015). Fused lasso additive model. *Journal of Computational and Graphical Statistics, in press* 1–37.

Spira, A., Beane, J. E., Shah, V., Steiling, K., Liu, G., Schembri, F., Gilman, S., Dumas, Y.-M., Calner, P., Sebastiani, P. et al. (2007). Airway epithelial gene expression in the diagnostic evaluation of smokers with suspect lung cancer. *Nature medicine* **13** 361–366.

Stewart, G. (1979). Pertubation bounds for the definite generalized eigenvalue problem. *Linear Algebra and Its Applications* **23** 69–85.

Stewart, G. and Sun, J. (1990). *Matrix Perturbation Theory*. Elsevier.

Tibshirani, R., Hastie, T., Narasimhan, B. and Chu, G. (2003). Class prediction by nearest shrunken centroids, with applications to DNA microarrays. *Statistical Science* **18** 104–117.

Vu, V. Q., Cho, J., Lei, J. and Rohe, K. (2013). Fantope projection and selection: A near-optimal convex relaxation of sparse PCA. In *Advances in Neural Information Processing Systems*.

Vu, V. Q. and Lei, J. (2013). Minimax sparse principal subspace estimation in high dimensions. *Annals of Statistics* **41** 2905–2947.

Wang, Z., Lu, H. and Liu, H. (2014). Tighten after relax: Minimax-optimal sparse PCA in polynomial time. In *Advances in Neural Information Processing Systems*.

Witten, D. M., Tibshirani, R. and Hastie, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics* **10** 515–534.

Yin, X. and Hilafu, H. (2015). Sequential sufficient dimension reduction for large p, small n problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **77** 879–892.

Yuan, X.-T. and Zhang, T. (2013). Truncated power method for sparse eigenvalue problems. *The Journal of Machine Learning Research* **14** 899–925.

Zhu, L., Miao, B. and Peng, H. (2006). On sliced inverse regression with high-dimensional covariates. *Journal of the American Statistical Association* **101** 630–643.

Zou, H., Hastie, T. and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics* **15** 265–286.

A Proof of Theorem 1

Next we state a series of lemmas that will be used in proving Theorem 1. The proofs for the technical lemmas are deferred to Appendix B. We start with some results from perturbation theory for eigenvalue and generalized eigenvalue problems (Golub and Van Loan, 2012).

Lemma 1. Let $\mathbf{J}$ and $\mathbf{J} + \mathbf{E}_\mathbf{J}$ be $d \times d$ symmetric matrices. Then, for all $k \in \{1, \ldots, d\}$,

$$\lambda_k(\mathbf{J}) + \lambda_{\min}(\mathbf{E}_\mathbf{J}) \leq \lambda_k(\mathbf{J} + \mathbf{E}_\mathbf{J}) \leq \lambda_k(\mathbf{J}) + \lambda_{\max}(\mathbf{E}_\mathbf{J}).$$

In the sequel, we state a result on the perturbed generalized eigenvalues for a symmetric-definite matrix pair $(\mathbf{J}, \mathbf{K})$ in the following lemma, which follows directly from Theorem 3.2 in Stewart (1979) and Theorem 8.7.3 in Golub and Van Loan (2012).

Lemma 2. Let $(\mathbf{J}, \mathbf{K})$ be a symmetric-definite matrix pair and let $(\mathbf{J}+\mathbf{E_J}, \mathbf{K}+\mathbf{E_K})$ be the perturbed matrix pair. Assume that $\mathbf{E_J}$ and $\mathbf{E_K}$ satisfy

$$\epsilon = \sqrt{\|\mathbf{E_J}\|_2^2 + \|\mathbf{E_K}\|_2^2} < \text{cr}(\mathbf{J}, \mathbf{K}),$$

where $\text{cr}(\mathbf{J}, \mathbf{K})$ is as defined in (9). Then, $(\mathbf{J}+\mathbf{E_J},\ \mathbf{K}+\mathbf{E_K})$ is a symmetric-definite matrix pair. Let $\lambda_k(\mathbf{J}+\mathbf{E_J},\ \mathbf{K}+\mathbf{E_K})$ be the kth generalized eigenvalue of the perturbed matrix pair. Then,

$$\frac{\lambda_k(\mathbf{J}, \mathbf{K}) \cdot \text{cr}(\mathbf{J}, \mathbf{K}) - \epsilon}{\text{cr}(\mathbf{J}, \mathbf{K}) + \epsilon \cdot \lambda_k(\mathbf{J}, \mathbf{K})} \leq \lambda_k(\mathbf{J}+\mathbf{E_J}, \mathbf{K}+\mathbf{E_K}) \leq \frac{\lambda_k(\mathbf{J}, \mathbf{K}) \cdot \text{cr}(\mathbf{J}, \mathbf{K}) + \epsilon}{\text{cr}(\mathbf{J}, \mathbf{K}) - \epsilon \cdot \lambda_k(\mathbf{J}, \mathbf{K})}.$$

Recall from Section 4 that $\mathbf{v}^*$ is the first generalized eigenvector of $(\mathbf{A}, \mathbf{B})$ with generalized eigenvalue λ_1, and that $V = \text{supp}(\mathbf{v}^*)$. For any given set F such that $V \subset F$, let $\lambda_k(F)$ and $\widehat{\lambda}_k(F)$ be the kth generalized eigenvalues of $(\mathbf{A}_F, \mathbf{B}_F)$ and $(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$, respectively. Provided Assumption 1 and by an application Lemma 2, we have

$$\frac{\widehat{\lambda}_2(F)}{\widehat{\lambda}_1(F)} \leq \gamma,$$

where $\gamma = (1+a)\lambda_2/[(1-a)\lambda_1]$ is as defined in (13). Let $\mathbf{y}(F) = \mathbf{v}(F)/\|\mathbf{v}(F)\|_2$ and $\mathbf{y}^* = \mathbf{v}^*/\|\mathbf{v}^*\|_2$ such that $\|\mathbf{y}(F)\|_2 = \|\mathbf{y}^*\|_2 = 1$. Next, we prove that $\mathbf{y}(F)$ is close to $\mathbf{y}^*$ if the set F contains the support of $\mathbf{y}^*$. The result follows from Theorem 4.3 in Stewart (1979).

Lemma 3. Let F be a set such that $V \subset F$ with $|F| = k' > s$ and let

$$\delta(F) = \sqrt{\|\mathbf{E}_{\mathbf{A},F}\|_2^2 + \|\mathbf{E}_{\mathbf{B},F}\|_2^2}.$$

Let

$$\chi(\lambda_1(F), \widehat{\lambda}_k(F)) = \frac{|\lambda_1(F) - \widehat{\lambda}_k(F)|}{\sqrt{1+\lambda_1(F)^2} \cdot \sqrt{1+\widehat{\lambda}_k(F)^2}}; \qquad \Delta\widehat{\lambda}(F) = \min_{k>1} \chi(\lambda_1(F), \widehat{\lambda}_k(F)) > 0.$$

If $\delta(F)/\Delta\widehat{\lambda}(F) < \text{cr}(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$, then

$$\frac{\|\mathbf{v}(F) - \mathbf{v}^*\|_2}{\|\mathbf{v}^*\|_2} \leq \frac{\delta(F)}{\Delta\widehat{\lambda}(F) \cdot \text{cr}(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)}.$$

This implies that

$$\|\mathbf{y}(F) - \mathbf{y}^*\|_2 \leq \frac{2}{\Delta\lambda \cdot \text{cr}(k')} \cdot \epsilon(k'),$$

where $\Delta\lambda$, $\text{cr}(k')$, and $\epsilon(k')$ are as defined in (12) and (10).

We now present a key lemma on measuring the progress of the gradient descent step. It requires an initial solution that is close enough to the optimal value in (15). With some abuse of notation, we indicate $\mathbf{y}(F)$ to be a k'-dimensional vector restricted to the set $F \subset \{1, \ldots, d\}$ with $|F| = k'$. Recall that $c > 0$ is some arbitrary small constant stated in Assumption 1 and c_{upper} is defined as $(1+c)/(1-c)$.

Lemma 4. Let $F \subset \{1, \ldots, d\}$ be some set with $|F| = k'$. Given any $\widetilde{\mathbf{v}}$ such that $\|\widetilde{\mathbf{v}}\|_2 = 1$ and $\widetilde{\mathbf{v}}^T \mathbf{y}(F) > 0$, let $\rho = \widetilde{\mathbf{v}}^T \widehat{\mathbf{A}}_F \widetilde{\mathbf{v}} / \widetilde{\mathbf{v}}^T \widehat{\mathbf{B}}_F \widetilde{\mathbf{v}}^2$, and let $\mathbf{v}' = \mathbf{C}_F \widetilde{\mathbf{v}} / \|\mathbf{C}_F \widetilde{\mathbf{v}}\|_2$, where

$$\mathbf{C} = \mathbf{I} + (\eta/\rho)(\widehat{\mathbf{A}} - \rho \widehat{\mathbf{B}})$$

and $\eta > 0$ is some positive constant. Let $\delta = 1 - \mathbf{y}(F)^T \widetilde{\mathbf{v}}$. If η is sufficiently small such that

$$\eta \lambda_{\max}(\mathbf{B}) < 1/(1+c),$$

and δ is sufficiently small such that

$$\delta \leq \min\left(\frac{1}{8 c_{\text{upper}} \kappa(\mathbf{B})}, \frac{1/\gamma - 1}{3 c_{\text{upper}} \kappa(\mathbf{B})}, \frac{1-\gamma}{30 \cdot (1+c) \cdot c_{\text{upper}}^2 \cdot \eta \cdot \lambda_{\max}(\mathbf{B}) \cdot \kappa^2(\mathbf{B}) \cdot [c_{\text{upper}} \kappa(\mathbf{B}) + \gamma]} \right),$$

then under Assumption 1, we have

$$\mathbf{y}(F)^T \mathbf{v}' \geq \mathbf{y}(F)^T \widetilde{\mathbf{v}} + \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot [1 - \mathbf{y}(F)^T \widetilde{\mathbf{v}}] \cdot \left[\frac{1-\gamma}{c_{\text{upper}} \kappa(\mathbf{B}) + \gamma} \right].$$

The following lemma characterizes the error introduced by the truncation step. It follows directly from Lemma 12 in Yuan and Zhang (2013).

Lemma 5. Consider $\mathbf{y}'$ with $F' = \text{supp}(\mathbf{y}')$ and $|F'| = \bar{k}$. Let F be the indices of $\mathbf{y}$ with the largest k absolute values, with $|F| = k$. If $\|\mathbf{y}'\|_2 = \|\mathbf{y}\|_2 = 1$, then

$$\begin{aligned} &|\text{Truncate}(\mathbf{y}, F)^T \mathbf{y}'| \\ &\quad \geq |\mathbf{y}^T \mathbf{y}'| - (\bar{k}/k)^{1/2} \min\left(\sqrt{1 - (\mathbf{y}^T \mathbf{y}')^2}, [1 + (\bar{k}/k)^{1/2}] \cdot [1 - (\mathbf{y}^T \mathbf{y}')^2] \right). \end{aligned}$$

The following lemma quantifies the progress of each iteration of Algorithm 1.

Lemma 6. Assume that $k > s$, where s is the cardinality of the support of $\mathbf{y}^* = \mathbf{v}^* / \|\mathbf{v}^*\|_2^2$, and k is the truncation parameter in Algorithm 1. Let $k' = 2k + s$ and let

$$\nu = \sqrt{1 + 2[(s/k)^{1/2} + s/k]} \cdot \sqrt{1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}} \kappa(\mathbf{B}) + \gamma} \right]}.$$

Under the same conditions in Lemma 4, we have

$$\sqrt{1 - |(\mathbf{y}^*)^T \widehat{\mathbf{v}}_t|} \leq \nu \sqrt{1 - |(\mathbf{y}^*)^T \mathbf{v}_{t-1}|} + \sqrt{10} \cdot \omega(k'),$$

In the following, we proceed with the proof of Theorem 1. Recall from Algorithm 1 that we define $\mathbf{v}_t = \widehat{\mathbf{v}}_t / \|\widehat{\mathbf{v}}_t\|_2$. Since $\|\mathbf{v}'_t\|_2 = 1$, and $\widehat{\mathbf{v}}_t$ is the truncated version of $\mathbf{v}'_t$, we have that $\|\widehat{\mathbf{v}}_t\|_2 \leq 1$. This implies that $|(\mathbf{y}^*)^T \mathbf{v}_t| \geq |(\mathbf{y}^*)^T \widehat{\mathbf{v}}_t|$. By Lemma 6, we have

$$\begin{aligned} \sqrt{1 - |(\mathbf{y}^*)^T \mathbf{v}_t|} &\leq \sqrt{1 - |(\mathbf{y}^*)^T \widehat{\mathbf{v}}_t|} \\ &\leq \nu \sqrt{1 - |(\mathbf{y}^*)^T \mathbf{v}_{t-1}|} + \sqrt{10} \cdot \omega(k'). \end{aligned} \tag{16}$$

By recursively applying Lemma 6, we have for all $t \geq 0$,

$$\sqrt{1 - |(\mathbf{y}^*)^T \mathbf{v}_t|} \leq \nu^t \sqrt{1 - |(\mathbf{y}^*)^T \mathbf{v}_0|} + \sqrt{10} \cdot \omega(k')/(1-\nu),$$

as desired.

B Proof of Technical Lemmas

B.1 Proof of Lemma 3

Proof. The first part of the lemma on the following inequality follows directly from Theorem 4.3 in Stewart (1979)

$$\frac{\|\mathbf{v}(F) - \mathbf{v}^*\|_2}{\|\mathbf{v}^*\|_2} \leq \frac{\delta(F)}{\Delta\widehat{\lambda} \cdot \mathrm{cr}(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)}.$$

We now prove the second part of the lemma.

Setting $\mathbf{y}(F) = \mathbf{v}(F)/\|\mathbf{v}(F)\|_2$ and $\mathbf{y}^* = \mathbf{v}^*/\|\mathbf{v}^*\|_2$ such that $\|\mathbf{y}(F)\|_2 = 1$ and $\|\mathbf{y}^*\|_2 = 1$, we have

$$\begin{aligned} \|\mathbf{y}(F) - \mathbf{y}^*\|_2 &\leq \left\| \frac{\mathbf{v}(F)}{\|\mathbf{v}(F)\|_2} - \frac{\mathbf{v}^*}{\|\mathbf{v}^*\|_2} \right\|_2 \\ &\leq \frac{1}{\|\mathbf{v}(F)\|_2 \cdot \|\mathbf{v}^*\|_2} \cdot \|\mathbf{v}(F) \cdot \|\mathbf{v}^*\|_2 - \mathbf{v}^* \cdot \|\mathbf{v}(F)\|_2\|_2 \\ &\leq \frac{2}{\|\mathbf{v}^*\|_2} \cdot \|\mathbf{v}(F) - \mathbf{v}^*\|_2 \\ &\leq 2\frac{\delta(F)}{\Delta\widehat{\lambda} \cdot \mathrm{cr}(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)} \end{aligned}$$

where the third inequality holds by adding and subtracting $\mathbf{v}(F) \cdot \|\mathbf{v}(F)\|_2$. By definition, $\delta(F) \leq \epsilon(k')$, $\Delta\widehat{\lambda} \geq \Delta\lambda$, and $\mathrm{cr}(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F) \geq \mathrm{cr}(k')$. Thus, we obtain

$$\|\mathbf{y}(F) - \mathbf{y}^*\|_2 \leq \frac{2}{\Delta\lambda \cdot \mathrm{cr}(k')} \cdot \epsilon(k').$$

□

B.2 Proof of Lemma 4

Proof. Recall that $F \subset \{1, \ldots, d\}$ is some set with cardinality $|F| = k'$. Also, recall that $\mathbf{y}(F)$ is proportional to the largest generalized eigenvector of $(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$. Throughout the proof, we write $\widehat{\kappa}$ to denote $\kappa(\widehat{\mathbf{B}}_F)$ for notational convenience. In addition, we use the notation $\|\mathbf{v}\|^2_{\widehat{\mathbf{B}}_F}$ to indicate $\mathbf{v}^T \widehat{\mathbf{B}}_F \mathbf{v}$.

Let $\boldsymbol{\xi}_j$ be the jth generalized eigenvector of $(\widehat{\mathbf{A}}_F, \widehat{\mathbf{B}}_F)$ corresponding to $\widehat{\lambda}_j(F)$ such that

$$\boldsymbol{\xi}_j^T \widehat{\mathbf{B}}_F \boldsymbol{\xi}_k = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k. \end{cases}$$

Assume that $\widetilde{\mathbf{v}} = \sum_{j=1}^{k'} \alpha_j \boldsymbol{\xi}_j$ and by definition we have $\mathbf{y}(F) = \boldsymbol{\xi}_1 / \|\boldsymbol{\xi}_1\|_2$. By assumption, we have $\mathbf{y}(F)^T \widetilde{\mathbf{v}} = 1 - \delta$. This implies that $\|\mathbf{y}(F) - \widetilde{\mathbf{v}}\|_2^2 = 2\delta$. Also, note that

$$\begin{aligned} \|\widetilde{\mathbf{v}} - \mathbf{y}(F)\|^2_{\widehat{\mathbf{B}}_F} &= \|\widetilde{\mathbf{v}} - \alpha_1 \boldsymbol{\xi}_1 - (\mathbf{y}(F) - \alpha_1 \boldsymbol{\xi}_1)\|^2_{\widehat{\mathbf{B}}_F} \\ &= \|\widetilde{\mathbf{v}} - \alpha_1 \boldsymbol{\xi}_1\|^2_{\widehat{\mathbf{B}}_F} + \|\mathbf{y}(F) - \alpha_1 \boldsymbol{\xi}_1\|^2_{\widehat{\mathbf{B}}_F} - 2[\mathbf{y}(F) - \alpha_1 \boldsymbol{\xi}_1]^T \widehat{\mathbf{B}}_F (\widetilde{\mathbf{v}} - \alpha_1 \boldsymbol{\xi}_1) \end{aligned}$$

Since $\mathbf{y}(F) - \alpha_1 \boldsymbol{\xi}_1$ is orthogonal to $\widetilde{\mathbf{v}} - \alpha_1 \boldsymbol{\xi}_1$ under the normalization of $\widehat{\mathbf{B}}_F$, we have

$$\sum_{j=2}^{k'} \alpha_j^2 = \|\widetilde{\mathbf{v}} - \alpha_1 \boldsymbol{\xi}_1\|^2_{\widehat{\mathbf{B}}_F} \leq \|\widetilde{\mathbf{v}} - \mathbf{y}(F)\|^2_{\widehat{\mathbf{B}}_F} \leq 2\lambda_{\max}(\widehat{\mathbf{B}}_F)\delta, \tag{17}$$

in which the last inequality holds by an application of Hölder's inequality and the fact that $\|\mathbf{y}(F) - \widetilde{\mathbf{v}}\|_2^2 = 2\delta$. Moreover, we have

$$\sum_{j=1}^{k'} \alpha_j^2 = \|\widetilde{\mathbf{v}}\|^2_{\widehat{\mathbf{B}}_F} \geq \lambda_{\max}(\widehat{\mathbf{B}}_F)/\widehat{\kappa} \qquad \text{and} \qquad \alpha_1^2 \geq \lambda_{\max}(\widehat{\mathbf{B}}_F)/\widehat{\kappa} - \sum_{j=2}^{k'} \alpha_j^2 \geq \frac{2\lambda_{\max}(\widehat{\mathbf{B}}_F)}{3\widehat{\kappa}}, \tag{18}$$

where the last inequality is obtained by (17) and the assumption that $\delta \leq 1/(8c_{\text{upper}}\kappa)$.

We also need a lower bound on $\|\mathbf{y}(F)\|_{\widehat{\mathbf{B}}_F}$. By the triangle inequality, we have

$$\begin{aligned} \|\mathbf{y}(F)\|_{\widehat{\mathbf{B}}_F} &\geq \|\widetilde{\mathbf{v}}\|_{\widehat{\mathbf{B}}_F} - \|\widetilde{\mathbf{v}} - \mathbf{y}(F)\|_{\widehat{\mathbf{B}}_F} \geq \sqrt{\sum_{j=1}^{k'} \alpha_j^2} - \sqrt{\lambda_{\max}(\widehat{\mathbf{B}}_F)} \cdot \|\widetilde{\mathbf{v}} - \mathbf{y}(F)\|_2 \\ &\geq \frac{1}{2}\sqrt{\sum_{j=1}^{k'} \alpha_j^2} + \frac{1}{2}\sqrt{\frac{\lambda_{\max}(\widehat{\mathbf{B}}_F)}{\widehat{\kappa}}} - \sqrt{2\lambda_{\max}(\widehat{\mathbf{B}}_F)\delta} \geq \frac{1}{2}\alpha_1, \end{aligned} \tag{19}$$

where the second inequality holds by the definition of $\|\widetilde{\mathbf{v}}\|_{\widehat{\mathbf{B}}_F}$ and an application of Hölder's inequality, the third inequality follows from (18), and the last inequality follows from the fact

that $1/2 \cdot \sqrt{\lambda_{\max}(\widehat{\mathbf{B}}_F)/\widehat{\kappa}} \geq \sqrt{2\lambda_{\max}(\widehat{\mathbf{B}}_F)\delta}$ under the assumption that $1/(8c_{\text{upper}}\kappa)$.

Lower and upper bounds for $[\widehat{\lambda}_1(F)-\rho]/\rho$: To obtain a lower bound for the quantity $\mathbf{y}(F)^T\mathbf{v}'$, we need both lower bound and upper bound for the quantity $[\widehat{\lambda}_1(F)-\rho]/\rho$. Recall that $\rho = \widetilde{\mathbf{v}}^T\widehat{\mathbf{A}}_F\widetilde{\mathbf{v}}/\widetilde{\mathbf{v}}^T\widehat{\mathbf{B}}_F\widetilde{\mathbf{v}}$. Using the fact that $\widetilde{\mathbf{v}}^T\widehat{\mathbf{A}}_F\widetilde{\mathbf{v}} = \sum_{j=1}^{k'}\alpha_j^2\widehat{\lambda}_j(F)$, we obtain

$$\frac{\widehat{\lambda}_1(F)-\rho}{\rho} = \frac{\sum_{j=1}^{k'}[\widehat{\lambda}_1(F)-\widehat{\lambda}_j(F)]\alpha_j^2}{\sum_{j=1}^{k'}\widehat{\lambda}_j(F)\alpha_j^2} \leq \frac{\widehat{\lambda}_1(F)\sum_{j=2}^{k'}\alpha_j^2}{\widehat{\lambda}_1(F)\alpha_1^2} \leq \frac{2\lambda_{\max}(\widehat{\mathbf{B}}_F)\delta}{\alpha_1^2} \leq 3\delta\widehat{\kappa}, \tag{20}$$

where the second to the last inequality holds by (17) and the last inequality holds by (18). We now establish a lower bound for $[\widehat{\lambda}_1(F)-\rho]/\rho$. First, we observe that

$$\delta \leq 2\delta - \delta^2 = (1-\delta)^2 + 1 - 2(1-\delta)\mathbf{y}(F)^T\widetilde{\mathbf{v}} = \|\widetilde{\mathbf{v}} - (1-\delta)\mathbf{y}(F)\|_2^2 \leq \|\widetilde{\mathbf{v}} - \alpha_1\boldsymbol{\xi}_1\|_2^2, \tag{21}$$

where the first equality follows from the fact that $\mathbf{y}(F)^T\widetilde{\mathbf{v}} = 1-\delta$, and the second inequality holds by the fact that $(1-\delta)\mathbf{y}(F)$ is the scalar projection of $\mathbf{y}(F)$ onto the vector $\boldsymbol{\xi}_1$. Thus, we have

$$\begin{aligned}
\frac{\widehat{\lambda}_1(F)-\rho}{\rho} &= \frac{\sum_{j=1}^{k'}[\widehat{\lambda}_1(F)-\widehat{\lambda}_j(F)]\alpha_j^2}{\sum_{j=1}^{k'}\widehat{\lambda}_j(F)\alpha_j^2} \geq \frac{[\widehat{\lambda}_1(F)-\widehat{\lambda}_2(F)]\sum_{j=2}^{k'}\alpha_j^2}{\widehat{\lambda}_1(F)\alpha_1^2 + \widehat{\lambda}_2(F)\sum_{j=2}^{k'}\alpha_j^2} \\
&= \frac{[\widehat{\lambda}_1(F)-\widehat{\lambda}_2(F)]\cdot\|\widetilde{\mathbf{v}}-\alpha_1\boldsymbol{\xi}_1\|_{\widehat{\mathbf{B}}_F}^2}{\widehat{\lambda}_1(F)\alpha_1^2 + \widehat{\lambda}_2(F)\cdot\|\widetilde{\mathbf{v}}-\alpha_1\boldsymbol{\xi}_1\|_{\widehat{\mathbf{B}}_F}^2} \geq \frac{(1-\gamma)\cdot[\lambda_{\max}(\widehat{\mathbf{B}}_F)/\widehat{\kappa}]\cdot\|\widetilde{\mathbf{v}}-\alpha_1\boldsymbol{\xi}_1\|_2^2}{\alpha_1^2 + \gamma\cdot[\lambda_{\max}(\widehat{\mathbf{B}}_F)/\widehat{\kappa}]\cdot\|\widetilde{\mathbf{v}}-\alpha_1\boldsymbol{\xi}_1\|_2^2} \\
&\geq \frac{(1-\gamma)\cdot\lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot\delta}{\alpha_1^2\cdot\widehat{\kappa} + \gamma\cdot\lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot\delta},
\end{aligned} \tag{22}$$

where the second to the last inequality holds by dividing the numerator and denominator by $\widehat{\lambda}_1(F)$ and using the upper bound $\widehat{\lambda}_2(F)/\widehat{\lambda}_1(F) \leq \gamma$, and the last inequality holds by (21).

Lower bound for $\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^{-1}$: In the sequel, we first establish an upper bound for $\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^2$. By the definition that $\rho = \widetilde{\mathbf{v}}^T\widehat{\mathbf{A}}_F\widetilde{\mathbf{v}}/\widetilde{\mathbf{v}}^T\widehat{\mathbf{B}}_F\widetilde{\mathbf{v}}$, we have

$$\widetilde{\mathbf{v}}^T\widehat{\mathbf{A}}_F\widetilde{\mathbf{v}} - \rho\widetilde{\mathbf{v}}^T\widehat{\mathbf{B}}_F\widetilde{\mathbf{v}} = 0.$$

Moreover, by the definition of $\widetilde{\mathbf{v}} = \sum_{j=1}^{k'}\alpha_j\boldsymbol{\xi}_j$ and the fact that $\widehat{\mathbf{A}}_F\boldsymbol{\xi}_j = \widehat{\lambda}_j(F)\widehat{\mathbf{B}}_F\boldsymbol{\xi}_j$, we have

$$\|(\widehat{\mathbf{A}}_F - \rho\widehat{\mathbf{B}}_F)\widetilde{\mathbf{v}}\|_2^2 = \left\|\sum_{j=1}^{k'}\alpha_j\widehat{\mathbf{A}}_F\boldsymbol{\xi}_j - \rho\sum_{j=1}^{k'}\alpha_j\widehat{\mathbf{B}}_F\boldsymbol{\xi}_j\right\|_2^2 = \left\|\sum_{j=1}^{k'}\alpha_j[\widehat{\lambda}_j(F)-\rho]\widehat{\mathbf{B}}_F\boldsymbol{\xi}_j\right\|_2^2. \tag{23}$$

Thus, by (23) and the fact that $\widetilde{\mathbf{v}}^T\widehat{\mathbf{A}}_F\widetilde{\mathbf{v}} - \rho\widetilde{\mathbf{v}}^T\widehat{\mathbf{B}}_F\widetilde{\mathbf{v}} = 0$, we obtain

$$\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^2 = \left\|\left[\mathbf{I} + \frac{\eta}{\rho}(\widehat{\mathbf{A}}_F - \rho\widehat{\mathbf{B}}_F)\right]\widetilde{\mathbf{v}}\right\|_2^2 = 1 + \left\|\sum_{j=1}^{k'}\alpha_j\cdot\left(\frac{\eta}{\rho}\right)\cdot[\widehat{\lambda}_j(F)-\rho]\cdot\widehat{\mathbf{B}}_F\boldsymbol{\xi}_j\right\|_2^2. \tag{24}$$

It remains to establish an upper bound for the second term in the above equation. Note that by the assumption that $\delta \le 1/(3\cdot c_{\text{upper}}\kappa)\cdot(1/\gamma - 1)$ and (20), we have

$$\widehat{\lambda}_2(F) \le \rho \le \widehat{\lambda}_1(F).$$

Moreover, since $\|\widetilde{\mathbf{v}}\|_2^2 = 1$, we have $\alpha_1^2 \le \lambda_{\max}(\widehat{\mathbf{B}}_F)$. Thus,

$$\begin{aligned}
&\left\|\sum_{j=1}^{k'}\alpha_j\cdot\left(\frac{\eta}{\rho}\right)\cdot[\widehat{\lambda}_j(F)-\rho]\cdot\widehat{\mathbf{B}}_F\boldsymbol{\xi}_j\right\|_2^2 \\
&\quad\le \alpha_1^2(\widehat{\lambda}_1(F)-\rho)^2\lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot(\eta/\rho)^2 + \lambda_{\max}(\widehat{\mathbf{B}}_F)\sum_{j=2}^{k'}\alpha_j^2\cdot(\eta/\rho)^2[\widehat{\lambda}_j(F)-\rho]^2 \\
&\quad\le \lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot(3\delta\widehat{\kappa})^2 + \lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot[\widehat{\lambda}_1(F)/\rho-1]^2\cdot\sum_{j=2}^{k'}\alpha_j^2 \\
&\quad\le \lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot(3\delta\widehat{\kappa})^2 + 2\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta\cdot(3\delta\widehat{\kappa})^2 \\
&\quad= 9\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2 + 18\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^3\cdot\widehat{\kappa}^2,
\end{aligned} \tag{25}$$

where the second inequality is from (20) and the third inequality follows from (17). Substituting (25) into (24), we have

$$\begin{aligned}
\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^2 &\le 1 + 9\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2 + 18\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^3\cdot\widehat{\kappa}^2 \\
&\le 1 + 12\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2,
\end{aligned} \tag{26}$$

where the last inequality follows from the fact that $2\delta \le 1/4$, which holds by the assumption that $\delta \le 1/(8c_{\text{upper}}\kappa)$. Meanwhile, note that the second term in the upper bound is less than one by the assumption $\delta \le 1/(8c_{\text{upper}}\kappa)$ and $\eta c_{\text{upper}}\lambda_{\max}(\mathbf{B}) < 1$. Hence, by invoking (26) nd the fact that $1/\sqrt{1+y} \ge 1 - y/2$ for $|y| < 1$, we have

$$\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^{-1} \ge 1 - 6\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2. \tag{27}$$

Lower bound for $\mathbf{y}(F)^T\mathbf{C}_F\widetilde{\mathbf{v}}$: We have

$$\begin{aligned}
\mathbf{y}(F)^T\mathbf{C}_F\widetilde{\mathbf{v}} &= \mathbf{y}(F)^T\widetilde{\mathbf{v}} + \frac{\eta}{\rho}\cdot\mathbf{y}(F)^T(\widehat{\mathbf{A}}_F - \rho\widehat{\mathbf{B}}_F)\widetilde{\mathbf{v}} \\
&= 1-\delta+\frac{\eta}{\rho}\cdot[\widehat{\lambda}_1(F)-\rho]\cdot\mathbf{y}(F)^T\widehat{\mathbf{B}}_F\widetilde{\mathbf{v}} \\
&= 1-\delta+\frac{\eta}{\rho}\cdot[\widehat{\lambda}_1(F)-\rho]\cdot\left(\alpha_1\cdot\frac{\boldsymbol{\xi}_1^T\widehat{\mathbf{B}}_F\boldsymbol{\xi}_1}{\|\boldsymbol{\xi}_1\|_2}\right) \\
&= 1-\delta+\eta\cdot\alpha_1\cdot\left[\frac{\widehat{\lambda}_1(F)-\rho}{\rho}\right]\cdot\|\mathbf{y}(F)\|_{\widehat{\mathbf{B}}_F} \\
&\geq 1-\delta+\frac{1}{2}\cdot\eta\cdot\alpha_1^2\cdot\left[\frac{(1-\gamma)\cdot\lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot\delta}{\alpha_1^2\cdot\widehat{\kappa}+\gamma\cdot\lambda_{\max}(\widehat{\mathbf{B}}_F)\cdot\delta}\right] \\
&\geq 1-\delta+\frac{1}{2}\cdot\eta\cdot\frac{\alpha_1^2\cdot(1-\gamma)\cdot\delta}{\widehat{\kappa}+\gamma} \\
&\geq 1-\delta+\frac{1}{3}\cdot\eta\cdot\lambda_{\min}(\widehat{\mathbf{B}}_F)\cdot\frac{(1-\gamma)\cdot\delta}{(\widehat{\kappa}+\gamma)},
\end{aligned} \tag{28}$$

where the first inequality follows from (19) and (22), the second inequality uses the fact that $\alpha_1^2 \leq \lambda_{\max}(\widehat{\mathbf{B}}_F)$, and the last inequality follows from (18).

Combining the results: We now establish a lower bound on $\mathbf{y}(F)^T\mathbf{v}'$. From (27) and (28), we have

$$\begin{aligned}
\mathbf{y}(F)^T\mathbf{v}' &= \mathbf{y}(F)^T\mathbf{C}_F\widetilde{\mathbf{v}}\cdot\|\mathbf{C}_F\widetilde{\mathbf{v}}\|_2^{-1} \\
&\geq \left(1-\delta+\frac{1}{3}\cdot\eta\cdot\lambda_{\min}(\widehat{\mathbf{B}}_F)\cdot\left[\frac{(1-\gamma)\cdot\delta}{(\widehat{\kappa}+\gamma)}\right]\right)\cdot\left(1-6\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2\right) \\
&\geq 1-\delta+\frac{1}{3}\cdot\eta\cdot\lambda_{\min}(\widehat{\mathbf{B}}_F)\cdot\left[\frac{(1-\gamma)\cdot\delta}{(\widehat{\kappa}+\gamma)}\right]-6\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2 \\
&\quad -2\cdot\widehat{\kappa}^2\cdot\eta^3\cdot\lambda_{\max}^3(\widehat{\mathbf{B}}_F)\cdot\delta^2\cdot\left[\frac{(1-\gamma)\cdot\delta}{(\widehat{\kappa}+\gamma)}\right] \\
&\geq 1-\delta+\frac{1}{3}\cdot\eta\cdot\lambda_{\min}(\widehat{\mathbf{B}}_F)\cdot\left[\frac{(1-\gamma)\cdot\delta}{(\widehat{\kappa}+\gamma)}\right]-6.25\cdot\lambda_{\max}^2(\widehat{\mathbf{B}}_F)\cdot\eta^2\cdot\delta^2\cdot\widehat{\kappa}^2 \\
&\geq 1-\delta+\frac{1}{8}\cdot\eta\cdot\lambda_{\min}(\widehat{\mathbf{B}}_F)\cdot\left[\frac{(1-\gamma)\delta}{(\widehat{\kappa}+\gamma)}\right],
\end{aligned} \tag{29}$$

in which the third inequality holds by the assumption that the step size η is sufficiently small such that $\eta\lambda_{\max}(\widehat{\mathbf{B}}_F) < 1$, and the last inequality holds under the condition that

$$\frac{1-\gamma}{(\widehat{\kappa}+\gamma)} \geq 30\eta\lambda_{\max}(\widehat{\mathbf{B}})\delta\widehat{\kappa}^2,$$

which is implied by the following inequality under Assumption 1

$$\delta \leq \frac{1-\gamma}{30 \cdot (1+c) \cdot c_{\text{upper}}^2 \cdot \eta \cdot \lambda_{\max}(\mathbf{B}) \cdot \kappa^2 \cdot (c_{\text{upper}}\kappa + \gamma)}.$$

By Assumption 1, we have

$$\mathbf{y}(F)^T \mathbf{v}' \geq 1 - \delta + \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot [1 - \mathbf{y}(F)^T \widetilde{\mathbf{v}}] \cdot \left(\frac{1-\gamma}{c_{\text{upper}}\kappa + \gamma} \right),$$

as desired.

□

B.3 Proof of Lemma 6

Proof. Recall that V is the support of $\mathbf{v}^*$, the population leading generalized vector and also $\mathbf{y}^* = \mathbf{v}^*/\|\mathbf{v}^*\|_2$. Let $F_{t-1} = \text{supp}(\mathbf{v}_{t-1})$, $F_t = \text{supp}(\mathbf{v}_t)$, and let $F = F_{t-1} \cup F_t \cup V$. Note that the cardinality of F is no more than $k' = 2k + s$, since $|F_t| = |F_{t-1}| = k$. Let

$$\mathbf{v}_t' = \mathbf{C}_F \mathbf{v}_{t-1} / \|\mathbf{C}_F \mathbf{v}_{t-1}\|_2,$$

where $\mathbf{C}_F$ is the submatrix of $\mathbf{C}_F$ restricted to the rows and columns indexed by F. We note that $\mathbf{v}_t'$ is equivalent to the one in Algorithm 1, since the elements of $\mathbf{v}_t'$ outside of the set F take value zero.

Applying Lemma 4 with the set F, we obtain

$$\mathbf{y}(F)^T \mathbf{v}_t' \geq \mathbf{y}(F)^T \mathbf{v}_{t-1} + \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot [1 - \mathbf{y}(F)^T \mathbf{v}_{t-1}] \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma} \right].$$

Subtracting both sides of the equation by one and rearranging the terms, we obtain

$$1 - \mathbf{y}(F)^T \mathbf{v}_t' \leq [1 - \mathbf{y}(F)^T \mathbf{v}_{t-1}] \cdot \left\{ 1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma} \right] \right\}. \tag{30}$$

This implies that

$$\|\mathbf{y}(F) - \mathbf{v}_t'\|_2 \leq \|\mathbf{y}(F) - \mathbf{v}_{t-1}\|_2 \cdot \sqrt{1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma} \right]}. \tag{31}$$

By Lemma 3, we have $\|\mathbf{y}(F) - \mathbf{y}\|_2 \leq \omega(k')$. By the triangle inequality, we have

$$\begin{aligned}
\|\mathbf{y} - \mathbf{v}_t'\|_2 &\leq \|\mathbf{y}(F) - \mathbf{v}_t'\|_2 + \|\mathbf{y}(F) - \mathbf{y}\|_2 \\
&\leq \|\mathbf{y}(F) - \mathbf{v}_{t-1}\|_2 \cdot \sqrt{1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma} \right]} + \omega(k') \qquad (32) \\
&\leq \|\mathbf{y} - \mathbf{v}_{t-1}\|_2 \cdot \sqrt{1 - \frac{1+c}{8} \cdot \eta \cdot \lambda_{\min}(\mathbf{B}) \cdot \left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B}) + \gamma} \right]} + 2\omega(k'),
\end{aligned}$$

where the second inequality follows from (31). This is equivalent to

$$\sqrt{1-|\mathbf{y}^T\mathbf{v}_t'|} \le \sqrt{1-|\mathbf{y}^T\mathbf{v}_{t-1}|}\cdot\sqrt{1-\frac{1+c}{8}\cdot\eta\cdot\lambda_{\min}(\mathbf{B})\cdot\left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B})+\gamma}\right]}+\sqrt{2}\cdot\omega(k'). \tag{33}$$

We define

$$\nu=\sqrt{1+2[(s/k)^{1/2}+s/k]}\cdot\sqrt{1-\frac{1+c}{8}\cdot\eta\cdot\lambda_{\min}(\mathbf{B})\cdot\left[\frac{1-\gamma}{c_{\text{upper}}\kappa(\mathbf{B})+\gamma}\right]}.$$

By Lemma 5 and picking $k > s$, we have

$$\begin{aligned}\sqrt{1-|\mathbf{y}^T\widehat{\mathbf{v}}_t|} &\le \sqrt{1-|\mathbf{y}^T\mathbf{v}_t'|+[(s/k)^{1/2}+s/k]\cdot[1-|\mathbf{y}^T\mathbf{v}_t'|^2]}\\ &\le \sqrt{1-|\mathbf{y}^T\mathbf{v}_t'|}\cdot\sqrt{1+[(s/k)^{1/2}+s/k]\cdot[1+|\mathbf{y}^T\mathbf{v}_t'|]}\\ &\le \sqrt{1-|\mathbf{y}^T\mathbf{v}_t'|}\cdot\sqrt{1+2[(s/k)^{1/2}+s/k]}\\ &\le \nu\sqrt{1-|\mathbf{y}^T\mathbf{v}_{t-1}|}+\sqrt{10}\cdot\omega(k'),\end{aligned} \tag{34}$$

where the third inequality holds using the fact that $|\mathbf{y}^T\mathbf{v}_t'| \le 1$, and the last inequality holds by (33). □